Aleatoric and Epistemic Uncertainty in Conformal Prediction

Yusuf Sale[1,2] YUSUF.SALE@IFI.LMU.DE
Alireza Javanmardi[1,2] ALIREZA.JAVANMARDI@IFI.LMU.DE
Eyke Hüllermeier[1,2,3] EYKE@IFI.LMU.DE
[1]*LMU Munich,* [2]*MCML,* [3]*German Research Centre for Artificial Intelligence (DFKI)*

Editors: Khuong An Nguyen, Zhiyuan Luo, Harris Papadopoulos, Tuwe Löfström, Lars Carlsson and Henrik Boström

Keywords: aleatoric, epistemic, uncertainty quantification, conformal prediction

1. Introduction

Recently, there has been a particular interest in distinguishing different *types* of uncertainty in supervised machine learning (ML) settings (Hüllermeier and Waegeman, 2021). Aleatoric uncertainty captures the inherent randomness in the data-generating process. As it represents variability that cannot be reduced even with more data, it is often referred to as *irreducible* uncertainty. In contrast, epistemic uncertainty arises from a lack of knowledge about the underlying data-generating process, which—in principle—can be reduced by acquiring additional data or improving the model itself (*viz. reducible* uncertainty). In parallel, interest in conformal prediction (CP)—both its theory and applications—has become equally vigorous. Conformal Prediction (Vovk et al., 2005) is a model-agnostic framework for uncertainty quantification that provides prediction sets or intervals with rigorous statistical coverage guarantees. Notably, CP is distribution-free and makes only the mild assumption of exchangeability. Under this assumption, it yields prediction intervals that contain the true label with a user-specified probability. Thus, CP is seen as a promising tool to quantify uncertainty. But how is it related to aleatoric and epistemic uncertainty? In particular, we first analyze how (estimates of) aleatoric and epistemic uncertainty enter into the construction of vanilla CP—that is, how noise and model error jointly shape the global threshold. We then review "uncertainty-aware" extensions that integrate these uncertainty estimates into the CP pipeline.

2. Does CP account for aleatoric and epistemic uncertainty?

For the sake of exposure, assume $Y = f(X) + \varepsilon$. Further, write $\delta(X) = f(X) - \hat{f}(X)$. Then, the absolute residual used in standard CP can be written as $R = |\varepsilon + \delta(X)|$. Now, define the conditional cdf $F_{R|x}(t) := P\left(|\varepsilon + \delta(x)| \leq t \,|\, X = x\right)$, where we assume $\varepsilon \sim P_{\varepsilon|X=x}$. Clearly, if $\delta(x) = 0$ (perfect knowledge) then $F_{R|x}$ is *only* driven by aleatoric uncertainty (i.e., noise). If $\varepsilon = 0$ (deterministic world) then $F_{R|x}(t) = 1\{t \geq |\delta(x)|\}$, i.e., a point mass at the epistemic error. Now, let P_X be the (unknown) marginal distribution of X. By the law of total probability, $F_R(t) = P(R \leq t) = \int_{\mathcal{X}} F_{R|x}(t)\, dP_X(x)$. Thus, the *global* distribution of residuals (which is used in essence for split CP) is the P_X-weighted mixture of all the *local* residual distributions. Consequently, the population $(1-\alpha)$-quantile is $q_{1-\alpha} = \inf\{t : F_R(t) \geq 1-\alpha\}$. Obviously, this yields $\int_{\mathcal{X}} F_{R|x}(q_{1-\alpha})\, dP_X(x) = 1-\alpha$. Because each integrand $F_R(t)$ depends on both ε (aleatoric) and $\delta(x)$ (epistemic), their influence enters here only through the same integral. There is no way to *disentangle* them. Therefore, the single threshold $q_{1-\alpha}$ can not tell us whether a large future residual comes from aleatoric or from epistemic uncertainty at that particular instance. So, aleatoric and epistemic uncertainty is reflected in a *global* manner instead of *instance-wise*.

Homogeneous noise and constant model error. Suppose both sources of uncertainty are uniform across x, namely, $\varepsilon \sim P_\varepsilon$ (same distribution everywhere), and $\delta(x) = \delta_0$ (constant model error). Then for every x, $F_{R|x}(t) = P(|\varepsilon + \delta_0| \leq t) \equiv F_R(t)$, so no averaging over P_X is needed. The split CP quantile $q_{1-\alpha}$ is chosen so that $F_R(q_{1-\alpha}) = 1-\alpha$, and hence $F_{R|x}(q_{1-\alpha}) = 1-\alpha$ for all x. In this setting, conformal intervals enjoy conditional coverage at each x, and their width correctly reflects aleatoric and epistemic uncertainty instance-wise.

Heterogeneous noise and constant model error. Now suppose the model error is still constant, i.e., $\delta(x) = \delta_0$, but the noise varies with x. Concretely, we assume $\varepsilon \sim P_{\varepsilon|X=x}$ with $Var[\varepsilon|X=x] = \sigma^2(x)$. Then the local residual cdf is $F_{R|x} = P(|\varepsilon+\delta_0| \leq t|X = x) = P(|\varepsilon| \leq t - |\delta_0|) = F_{\varepsilon|x}(t - |\delta_0|)$, which clearly depends on $\sigma(x)$. However, the vanilla CP threshold is a single constant. The resulting interval has constant width, so it can not adapt to regions of higher or lower noise. To correct for varying noise, define an estimate $\hat{\sigma}(x) \approx \sigma(x)$ and define the normalized score (Papadopoulos et al., 2008; Lei et al., 2018) $U = (Y - \hat{f}(X))/\hat{\sigma}(X)$. Let $F_{U|x}(u) = P(|Y - \hat{f}(x)| \leq u\hat{\sigma}(x)|X = x) = F_{\varepsilon|x}(u\hat{\sigma}(x) - |\delta_0|)$. Crucially, if $\hat{\sigma}(x) = \sigma(x)$, then $F_{U|x}(u) = F_{\varepsilon|x}(\sigma(x)u - |\delta_0|)$ depends on x only through the shift δ_0, not through the scale. Under mild regularity, one can show that these $F_{U|x}$ are identical in shape across x. Hence the global cdf $F_U(u) = \int F_{U|x}(x)dP_X(x)$ coincides with each local $F_{U|x}(u)$. We then choose $T = \inf\{u : F_U(u) \geq 1-\alpha\}$. The resulting conformal interval has width $2\hat{\sigma}(x)T$ that scales with the local noise level. In contrast, conformalized quantile regression (CQR) first models the (local) noise via conditional quantiles and then adds a global slack to ensure validity (Romano et al., 2019). Thus, by normalizing residuals, CP can reflect aleatoric uncertainty per instance.

General setting. In the fully general setting—where both the noise $\sigma(x)$ and the model error $\delta(x)$ vary arbitrarily with x—one seeks a conformal score that adapts *both* aleatoric and epistemic uncertainty per instance. Uncertainty-Aware CQR (UCQR) replaces that single slack with an x-dependent adjustment derived from an ensemble's spread, so intervals widen precisely where model error is high (Rossellini et al., 2024). EPICSCORE takes this further by fitting a Bayesian predictive distribution of the nonconformity scores itself and using its local quantiles to define the interval at each x (Cabezas et al., 2025).

Outlook. A natural next step is to forge a more explicit dialogue between the "two-uncertainties" paradigm in ML—where aleatoric and epistemic uncertainty are formally disentangled—and the rich, distribution-free guarantees of CP. While we discussed these ideas in the context of regression, the same principles also apply to classification, as elaborated by Javanmardi et al. (2025). More generally, the notion of "instance-wise" uncertainty quantification is closely connected to the concept of "conditional coverage" in CP, which is certainly worth further exploration. In this light, one may ask whether CP ought to explicitly disentangle and adjust for aleatoric versus epistemic uncertainty.

Acknowledgments

Alireza Javanmardi was supported by the Klaus Tschira Stiftung. Yusuf Sale is supported by the DAAD program Konrad Zuse Schools of Excellence in Artificial Intelligence, sponsored by the Federal Ministry of Education and Research.

References

Luben Miguel Cruz Cabezas, Vagner Silva Santos, Thiago Ramos, and Rafael Izbicki. Epistemic uncertainty in conformal scores: A unified approach. In *The 41st Conference on Uncertainty in Artificial Intelligence*, 2025.

Eyke Hüllermeier and Willem Waegeman. Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods. *Machine learning*, 110(3):457–506, 2021.

Alireza Javanmardi, Soroush H Zargarbashi, Santo MAR Thies, Willem Waegeman, Aleksandar Bojchevski, and Eyke Hüllermeier. Optimal conformal prediction under epistemic uncertainty. *arXiv preprint arXiv:2505.19033*, 2025.

Jing Lei, Max G'Sell, Alessandro Rinaldo, Ryan J Tibshirani, and Larry Wasserman. Distribution-free predictive inference for regression. *Journal of the American Statistical Association*, 113(523):1094–1111, 2018.

Harris Papadopoulos, Alex Gammerman, and Volodya Vovk. Normalized nonconformity measures for regression conformal prediction. In *Proceedings of the IASTED International Conference on Artificial Intelligence and Applications (AIA 2008)*, pages 64–69, 2008.

Yaniv Romano, Evan Patterson, and Emmanuel Candes. Conformalized quantile regression. *Advances in neural information processing systems*, 32, 2019.

Raphael Rossellini, Rina Foygel Barber, and Rebecca Willett. Integrating uncertainty awareness into conformalized quantile regression. In *International Conference on Artificial Intelligence and Statistics*, pages 1540–1548. PMLR, 2024.

Vladimir Vovk, Alexander Gammerman, and Glenn Shafer. *Algorithmic learning in a random world*, volume 29. Springer, 2005.